

# Towards a Rigorous Basis for Specific Operations Risk Assessment of UAS

Ewen Denney, and Ganesh Pai  
*Intelligent Systems Division*  
 SGT / NASA Ames Research Center  
 Moffett Field, CA 94035, USA  
 {ewen.denney, ganesh.pai}@nasa.gov

Marcus Johnson  
*Aviation Systems Division*  
 NASA Ames Research Center  
 Moffett Field, CA 94035, USA  
 marcus.johnson@nasa.gov

**Abstract**—The Specific Operations Risk Assessment (SORA) guidance represents the consensus of various national aviation authorities on a common process to identify, qualitatively assess, and manage the safety risk posed by unmanned aircraft systems (UAS), when preparing the safety case required for regulatory approval to conduct certain types of operations. As such, it can be considered a de facto standard, being increasingly adopted by various relevant stakeholders. This paper first gives an overview of the SORA process and associated methods, identifying a number of inconsistencies in risk identification and assessment, also discussing plausible strategies to close the associated gaps. Then, we give a well-founded basis for the applicable concepts, such as barrier integrity, assurance, and robustness, following which we present a preliminary and simple probabilistic formalization of the underpinning barrier-based safety model. We illustrate our overall approach through a worked example, also discussing how a Bayesian framework can facilitate extending and enhancing our initial formalization. We conclude with a discussion of the opportunities afforded by our approach, such as a well-founded basis for barrier selection, whilst addressing the associated challenges. The main objective of this work is to complement the current SORA guidance through a principled, mathematically-based approach to risk assessment, particularly when it is applied to higher-risk operational concepts that warrant greater rigor in safety assessment and assurance.

**Index Terms**—Bayesian analysis, Risk assessment, Risk quantification, Safety architecture, Safety case, Unmanned aircraft systems.

## I. INTRODUCTION

*Specific Operations Risk Assessment* (SORA) [1] is both a methodology and guidance being promulgated by the *Joint Authorities for Rulemaking on Unmanned Systems* (JARUS)—an international consortium of national aviation authorities (NAAs) and regional aviation safety organizations—to support applications for authorization to operate unmanned aircraft systems (UAS) in civil airspace. At its core, SORA<sup>1</sup> is an approach to safety risk management (SRM), providing an uncomplicated and convenient means to qualitatively evaluate the safety risk associated with so-called *specific category* UAS operations, and determine whether those risks have been reduced to an acceptable level. Specific category operations are those whose level of safety risk falls between that posed by operations in the *open category* and the *certified category*. The

former constitutes very low risk operations, while the latter typically require type, airworthiness, and operator certification, as well as flight crew licensing [2].

In its current form, i.e., version 1.0 [1], SORA is based on a *barrier model* of safety which can be represented using *bow tie diagrams* (BTDs), a graphical elaboration of safety-relevant scenarios and the suite of related SRM measures (termed as *barriers*). Exemplifying a flexible and risk-based approach to trade off safety-relevant considerations—e.g., technical airworthiness, equipment and operator performance and capabilities, and operating rules, restrictions, and procedures—it effectively provides a basis for the operational safety case, so that the safety measures employed are proportional to the risk posed by the particular operational concept. An upcoming SORA revision<sup>2</sup> (i.e., version 2.0) has refocused on the application of the core methodology and the underlying process, though it continues to be implicitly based on BTDs. Section II gives a background on BTDs, SORA, and their interrelationship.

This paper adopts the position that the current, largely qualitative approach to (both versions of the) SORA can benefit from the rigor of a mathematical basis, especially when it is being applied to higher-risk operational concepts, e.g., beyond visual line of sight (BVLOS) overflight of urban, populated areas. Section III motivates our position based on gaps and inconsistencies we have identified in the SORA, and also discusses potential resolutions.

Then, in Section IV, we develop well-founded notions for various SORA concepts including *barrier integrity*, *assurance*, *correction factor*, *integrity level*, *assurance level*, *robustness*, and *robustness level*. Thereafter, we advance a preliminary, mathematically-based risk assessment approach, where we build a probabilistic model of the safety-relevant scenarios described using BTDs (e.g., those considered in the SORA guidance). We also discuss extending our model to a notion of *safety architecture* (SA) [3], itself an extension to BTDs.

We have implemented this model in our assurance case tool, AdvOCATE [4], which we used in creating a UAS safety case that successfully underwent regulatory scrutiny [5], to facilitate obtaining operational flight approval in the context of

<sup>1</sup>Henceforth, we will use the acronym in reference to the overall approach, the specific guidance document, as well as the constituent processes and methods, qualifying our usage when the context is unclear.

<sup>2</sup>Presently, drafts of SORA v2.0 and several of its annexes are undergoing a so-called *JARUS External Consultation*, where the documents undergo public review and feedback, prior to final publication and release. Available at: [http://jarus-rpas.org/sites/jarus-rpas.org/files/press\\_release\\_no\\_24.pdf](http://jarus-rpas.org/sites/jarus-rpas.org/files/press_release_no_24.pdf)

UAS Traffic Management (UTM) flight testing [6]. Based on this safety case, we present a simplified example to illustrate the feasibility and utility of the formalization as a well-founded implementation that conforms to the intent of the SORA.

Subsequently, Section V discusses extending and enhancing our preliminary model, by *i*) formalizing barrier integrity quantification, and capturing the role of *escalation factors* (EFs)—described next in Section II-A—that affect barrier performance, and *ii*) using a Bayesian probabilistic framework, to capture relations between barriers, and to quantify risk reduction as a query on the joint distribution of the associated random variables (RVs).

Then, Section VI discusses how our approach not only resolves the identified gaps (see Section III), but also provides opportunities for complementing SORA, e.g., updating the risk assessment based on operational data, including additional or related risk models, and providing a well-founded basis for selecting barriers and for specifying barrier integrities based on the desired safety targets. Finally, we consider the challenges associated with our approach, discuss potential solutions, and conclude highlighting avenues for future work.

## II. BACKGROUND

### A. Bow Tie Diagrams

*Bow tie diagrams* (BTDs), which realize a *barrier model* of safety, provide a graphical means to visualize and assess the risk scenarios associated with a given hazard [7]. Fig. 1 shows the main elements of a BTD, shown in graphical notation as implemented in our tool, AdvoCATE [4], [8] (see Fig. 3 for a more generic representation).

A *hazard* is a controlled activity, condition, or entity that reflects a normal, often desirable, aspect of the concept of operations (CONOPS), e.g., UAS operations near an airport. A *top event* is an undesired system state where control over the hazard is lost, e.g., a loss of separation from other aircraft. A *threat* is a possible direct cause or source of a top event, e.g., unanticipated course deviation. A *consequence* represents the dangerous outcome or loss state that results when a top event cannot be contained after it has occurred, e.g., a midair collision (MAC) with fatalities.

To manage threats and top events, we employ *controls*—any process, function, device, practice, or other action that modifies safety risk. A *barrier* is a collection or *system* of controls working, similarly, to modify (reduce) safety risk. For instance, a surveillance system represents a technical barrier comprising controls such as detection and tracking, classification, alerting and advisory, etc. *Escalation factors* (EFs) are weaknesses, vulnerabilities, threats, or operational conditions that can compromise, defeat, or otherwise degrade control effectiveness, e.g., electromagnetic interference. *Escalation factor barriers* (EFBs), analogous to barriers, are a *secondary* system of controls used to manage, reduce, or modify the impact of EFs.

*Prevention* barriers (controls) represent mitigation measures that target reducing the probability of the top event, and are shown preceding the top event. *Recovery* barriers (controls) contribute to reducing the probability and/or the magnitude

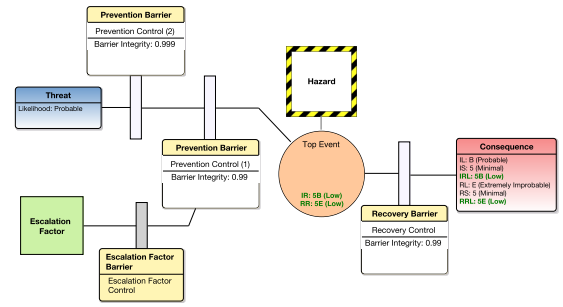


Fig. 1. Graphical representation of BTD elements as implemented in our safety assurance case tool, AdvoCATE.

of the severity of the consequence after a top event has occurred, and are shown following a top event. The visual ordering of barriers and controls loosely corresponds to the temporal order in which they may be invoked. However, BTDs (intentionally) abstract from specifying the exact ordering and organization of barriers (and their constituent controls) between successive events since those correspond to design decisions made typically after determining, at this abstract level, that the barriers being deployed can provide sufficient safety risk reduction.

Threats are assumed to occur independently, i.e., there is a possibility for threats to occur simultaneously and, therefore, they are not disjoint. Depending on the CONOPS, consequences may or may not be disjoint events. With this interpretation, threats, top events and consequences can be ascribed an *initial* and a *residual risk level*, computed as a combination of their (initial/residual) probabilities of occurrence and severity. Barriers and controls are each ascribed a measure of *integrity*, relating to the probability that barriers are (not) breached in a dangerous manner.

Although the BTD elements in SORA v1.0 are conceptually identical to those given here, SORA introduces minor terminological differences along with a few additional concepts (described next, in Section II-B). As mentioned earlier, SORA v2.0 is implicitly based on BTDs, referring to the more general, model-agnostic, terms of *mitigation* and *operational safety objective* (OSO), rather than harm and threat barrier, respectively. However, it retains the concepts of integrity, assurance, and robustness, whilst referring to a so-called *holistic risk model*. In fact, this is the generic BTD that underpins the methodology developed in SORA v1.0 [1].

### B. Specific Operations Risk Assessment

We mainly describe SORA v1.0 and, except for some simplifying modifications to the process steps, the core approach remains largely unchanged in SORA v2.0. Since the final revision of the latter is yet to be made publicly available, we do not give more details on it here. The SORA process starts with *risk modeling*, is followed by a *risk assessment*, and culminates with recommendations on the mitigation measures to be used for SRM. Fig. 2 shows the different tasks across these activities and the associated data flow, based on our understanding of the process.

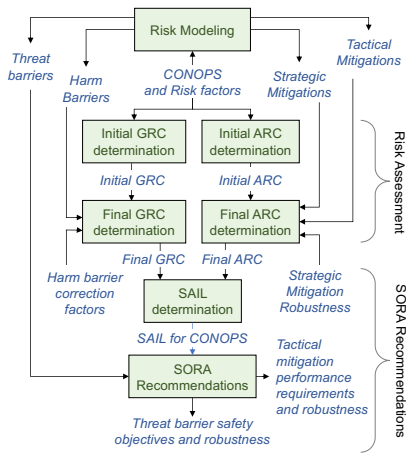


Fig. 2. SORA data flow starts with risk modeling and produces recommendations on risk mitigation barrier robustness, and the associated safety objectives.

1) *Risk Modeling*: Risk modeling, which is based on BTDS, focuses on a nonspecific operational hazard (*UAS operation out of control*) and three broad consequences—namely fatalities to non-participating third parties on the ground, in the air, and damage to critical infrastructure—whilst identifying generic threats, e.g., *technical issue with the UAS*, *adverse operating conditions*, etc. The focus, however, is mainly on the former two, i.e., the ground collision consequence (GCC), and the air collision consequence (ACC).

The results of risk modeling also include candidate *harm* barriers to mitigate consequence risk, and *threat* barriers to manage the identified threats.<sup>3</sup> An orthogonal organization of barriers that applies mainly to the ACC consists of so-called *strategic* and *tactical* mitigations. The former constitute barriers used *prior to* flight operations, while the latter are employed *during* flight. Depending on when a strategic or tactical mitigation is invoked, they can function in either a prevention or a recovery role.

2) *Risk Assessment*: The risk assessment activity comprises several steps that build upon each other to provide a qualitative risk assessment for the GCC and the ACC. Fig. 2 shows these steps as tasks that follow the initial risk modeling task, and which precede the SORA recommendations task.

a) *Initial Ground and Air Collision Risk Determination*: These activities start from the characteristics of the CONOPS that contribute to each consequence, e.g., aircraft dimensions, kinetic energy, operating area type, and population density are amongst the factors affecting GCC risk. Similarly, ACC risk depends on the operating airspace type and class, operating altitude, encounter rate, airspace geometry and structure, and aircraft dynamics. SORA uses these characteristics to classify and order operational scenarios in terms of their relative *unmitigated* risk, i.e., the risk posed without taking into account any risk reducing mitigation measures.

For the GCC, 8 broad operational scenarios have been identified, each assigned a *ground risk class* (GRC)—a relative

risk ranking from ‘1’ (lowest risk) to ‘10’ (highest risk)—whose intent is to capture the unmitigated risk that a person on the ground is struck by an out of control unmanned aircraft (UA). SORA currently focuses on operations whose GRC is no greater than ‘7’. Analogously for the ACC, 12 disjoint operational scenarios have been developed in the form of so-called *airspace encounter categories* (AECs), i.e., by classifying the operating airspace based on the intended operating altitude, airspace type, and airspace class. Each AEC is then associated with a perceived level of collision risk, expressed as an assignment to an *air risk class* (ARC), a relative risk ranking from ‘1’ (lowest risk) to ‘4’ (highest risk). Each ARC relates to the rate at which a UA would encounter a conventionally piloted aircraft (CPA)—i.e., aircraft with onboard human pilots—in that AEC.

Thus, ascribing a GRC and ARC provides a qualitative determination of the initial level of risk, in the context of the two main consequences of interest for a particular CONOPS. The SORA then describes simple procedures to refine this initial risk assessment into a final risk assessment.

b) *Final GRC Determination*: The core idea is to modify the initial GRC based on *correction factors* associated with the harm barriers. The correction factor for a given harm barrier is a numerical value selected from  $\{-4, -2, -1, 0, 1\}$ , indicating the (qualitative) extent of risk modification provided. Thus, ‘-4’ represents the largest amount of risk reduction, ‘0’ indicates no risk modification, while ‘1’ is a risk increase.

The correction factor value used depends on barrier *robustness level*, which the SORA qualifies as a combination of *level of integrity* (i.e., the ‘safety gain provided’), and *level of assurance* (i.e., the ‘proof that the claimed gain has been achieved’). Each of the three attributes are qualitatively gauged on a three-valued ordinal scale— $\{low, medium, high\}$ —and a mapping is specified that relates every combination of integrity and assurance level to robustness level. As a heuristic, for a given integrity level, robustness level is directly proportional to the assurance level.

The more robust a harm barrier is, the greater the amount of risk reduction it is purported to provide, and the larger its correction factor. The final risk assessment is then given as the sum of the GRC and the correction factors for all the harm barriers used. For example, if the initial GRC is ‘7’, say, and two harm barriers are used whose robustness level is such that their respective correction factors are -2 and -1, then the final GRC is given as  $7 + (-2) + (-1) = 4$ .

c) *Final ARC Determination*: Refining the initial ARC into the final ARC depends on applying the identified strategic mitigations and their associated robustness. The effect brought about by using these barriers is a reduction in the encounter rate, which is achieved by providing a) temporal and spatial operational restrictions that reduce exposure and air traffic density, and b) procedures and rules that restrict aircraft dynamics, as well as airspace structure and geometry.

In general, the greater the robustness is (or is shown to be) for a strategic mitigation, the greater the reduction claimed in the ARC. Then, to determine the final ARC, the SORA uses

<sup>3</sup>Harm and threat barriers are SORA v1.0 terminology for recovery and prevention barriers, respectively.

a qualitative process that takes into account: *i)* a qualitative robustness level for the strategic mitigations actually used, *ii)* a qualitative measure of the extent to which those mitigations modify the contribution of specific risk factors for the ACC, e.g., encounter rate, *iii)* verification of strategic mitigation robustness, and *iv)* validation of the mitigations used.

3) *Mitigation Recommendations*: The recommendation element of the SORA gives a mapping from the combination of the final ARC and GRC to specific threat/prevention barriers, via a so-called *Specific Assurance and Integrity Level* (SAIL), shown in Fig. 2 as the task *SAIL determination*. This is a qualitatively determined rank on a six-level ordinal scale from ‘I’ to ‘VI’ that corresponds to particular threat barriers (in SORA v1.0), or OSOs (in SORA v2.0) to be used, and the robustness they ought to exhibit.

SORA v1.0 associates each GRC and ARC, individually, with a particular SAIL and, in general, the latter increases proportionally with the former. Then, it gives the SAIL for a given CONOPS as the greater of the SAILs for the respective GRC and the ARC. SORA v2.0 simplifies this step, providing a single SAIL for every pair of GRC and ARC. Additionally, based on the final ARC and its corresponding SAIL, the SORA establishes so-called *tactical mitigation performance requirements* (TMPR)—classified as  $\{low, medium, high\}$ —representing a qualitative level of risk reduction to be achieved by the tactical mitigations to reduce the residual risk of the ACC. The final step is barrier robustness verification, which is, again, gauged qualitatively.

### III. MOTIVATING A MATHEMATICAL FOUNDATION

The overarching intent of SORA is to provide a (qualitative) level of confidence that a given UAS operation remains safely controlled. As described in the preceding section, upon applying the SORA to a CONOPS, one should expect to arrive at recommendations on barriers (or safety objectives) and their robustness, via a SAIL determination. Then by showing that the applicable safety objectives have been achieved, a sufficient (presumably *high*) level of confidence ought to be available that there is (acceptably) *low* GCC or ACC risk.

Thus, being practically inspired, and relatively uncomplicated in its approach to risk assessment and safety assurance, SORA shows promise of usefulness and is, therefore, attractive to adopt. Nevertheless, it presents a number of open issues and inconsistencies (three of which we describe next) that we believe ought to be resolved, and that motivate the advantage of mathematically-based foundations as a quantitative means for validating (certain) SORA assumptions. Later (Section IV), we will discuss how our formal model addresses these inconsistencies and open issues.

*i) The initial GRC of some of the identified ground collision scenarios are inconsistent with the corresponding BTDs.*

In particular, the initial GRC for the eight GCC scenarios identified<sup>4</sup>, relies on the following risk factors: *i)* the *expected kinetic energy* transferred in a collision given the dimensions

of the UA involved, *ii)* the *operation type*, i.e., within visual line of sight (VLOS) or BVLOS, *iii)* the *operating area type* (controlled or uncontrolled), and *iv)* *population density* in the area overflown, i.e., sparse, populated, etc. Given these, the scenarios involving BVLOS flight each have a higher initial GRC than the corresponding VLOS counterparts, with all other risk factors staying the same. For example, for a given UA, the initial GRC for ‘BVLOS operations over a gathering of people’ (‘8’) is greater than that of the corresponding VLOS scenario (‘7’). However, for this to be the case, BVLOS flight must inherently introduce a greater unmitigated probability of loss of UA control, when all other risk factors, including kinetic energy, remain unchanged in both scenarios.

So far as we can gauge, this is not evident from the generic BTD-based risk model underlying SORA—where the ACC and GCC are considered to be disjoint—thereby presenting an inconsistency. Moreover, the specific conditions or rationale for the same have not been clarified. We submit that BVLOS flight as a risk factor is largely relevant for the ACC rather than the GCC. We can eliminate this inconsistency by formally and explicitly modeling the, otherwise qualitative, relation between the GRCs, (initial unmitigated) GCC risk, and the contributing risk factors. As we will see later (Section IV), a mathematically founded basis for SORA tightly couples consequences and their risk factors, potentially allowing GRCs to be generated.

*ii) Modifying initial GRCs based on harm barrier correction factors can lead to an inconsistent risk assessment.*

First, both initial and final GRC, as well as correction factors (for a given robustness), can be viewed as a measurement on an ordinal scale [9]. That is, they mainly support an inference on relative risk (for GRCs), or relative risk modification (for correction factors), with no notion of a *degree of difference*. Indeed, without a mapping from GRCs to (a range of) absolute risk values, one can only meaningfully reason about a scenario posing greater or lesser risk than another scenario, rather than the amount to which the risk is greater or lesser. Similar constraints apply to correction factors of different harm barriers.

Next, not only are arithmetic transformations (summation in this case) inadmissible for ordinal scales, but also to apply them to correction factors and GRCs, there must be an equivalence between the underlying measurement scales, i.e., a correction factor of ‘1’ and GRC of ‘1’, say, must represent equivalent risk levels. So far as we are aware, SORA does not give a basis for this (assumption of) equivalence. Moreover the mapping from harm barrier correction factors to the actual level of risk modification provided, appears to be ad hoc.

We now give an example to illustrate the kind of problems that can arise when modifying GRCs using correction factors. First, we assume a (plausible) mapping from GRCs to the probability component of absolute risk values as: ‘5’  $\mapsto (10^{-3}, 10^{-2}]$ , ‘6’  $\mapsto (10^{-4}, 10^{-3}]$ , and ‘7’  $\mapsto (10^{-5}, 10^{-4}]$ . The interpretation is that a GRC of ‘5’, say, reflects a GCC probability  $> 10^{-3}$ , but  $\leq 10^{-2}$ , and so on. Now consider two different operating scenarios such that the GCC probability is slightly  $> 10^{-5}$  in the first scenario, and  $10^{-4}$  in the second,

<sup>4</sup>Figure 8 in SORA v1.0 [1], and Table 2 in the draft SORA v2.0.



so that they both have the same initial GRC of ‘7’ as per the preceding mapping. Further, assume that a harm barrier with a *medium* level of robustness is used, such that its correction factor is  $-2$ , say, corresponding to an absolute probability reduction of  $0.1$ , say. By applying the SORA correction factor arithmetic, the final GRC is  $7 + (-2) = 5$ , i.e., a GCC probability in the interval  $(10^{-3}, 10^{-2}]$ . However, in absolute terms, the actual probability reduction for the first scenario is slightly  $> 10^{-4}$ , i.e., GRC ‘6’, while for the second scenario, it is  $10^{-3}$ , i.e., also GRC ‘6’. Similar inconsistencies arise when considering the severity component of risk.

In the general case, depending on the mapping and the bounds chosen, such reasoning could impart a false assurance of risk reduction especially in high-risk scenarios. That, in turn, could adversely affect both the recommendations on the threat barriers to be used, and stakeholders’ assessment of the acceptability of residual risk. One possible resolution is to provide a mapping from GRCs and correction factors to the respective absolute values for risk and risk modification. However, this may not be generalizable across all potential stakeholders and operations. Another resolution is to associate GRCs and correction factors with descriptive or alphanumeric labels (e.g., ‘GRC-5’ instead of ‘5’, say) to preclude numerical manipulation. In fact, the SORA does exactly this for both initial and final ARCs, correctly regarding them as labels rather than numerical values.

iii) *There is limited explanation and validation of the basis for the relation between barrier robustness and the modification of the initial GRC or ARC.*

From the preceding discussion, evidently, it is problematic to apply harm barrier correction factors to numerically modify initial GRCs. As such, changing the value of correction factors based on harm barrier robustness, and, subsequently, using them to modify the initial GRC does not eliminate the induced inconsistencies. Additionally, as indicated earlier, the SORA approach to determine and modify the initial ARCs is different from that used for the GRCs.

To our knowledge, the JARUS consortium is continually refining the SORA, also developing the rationale for *a*) how AECs are established, given air collision risk factors, *b*) relating AECs to ARCs, *c*) the amount of risk modification a strategic mitigation affords, given its robustness, and *d*) the extent of risk modification that single or multiple strategic mitigations provide, in determining the final ARC. SORA v2.0 and its annexes, when published in their final form, may provide the required rationale for both strategic and tactical mitigations. Though we do not elaborate the specific details from the draft SORA v2.0, the overall approach—as we understand it—is a process-based qualitative determination of the degree to which strategic mitigations modify the initial ARC, along with prescriptive recommendations on the robustness that must be demonstrated for the modification to be considered valid.

Intuitively, however, it is reasonable to relate the robustness of a harm barrier or a strategic mitigation with the level of risk modification provided for the GCC and the ACC. To characterize the nature of this relation, a formal definition

of robustness would be required, together with a mapping to both the severity and probability component of GCC or ACC risk, itself expressed as an absolute (rather than relative) risk value. As we will see next, our formal basis facilitates such a mapping.

#### IV. PRELIMINARY FORMALIZATION

In this section, first we elaborate various barrier properties, giving them a formal basis and showing how they are consistent with the corresponding SORA characterizations. Then, we develop an objective, mathematical basis for SORA based on BTDs, explaining how they address the SORA inconsistencies and open issues identified in Section III.

##### A. Foundations for Barrier Properties

1) *Reliability*: We first consider barrier *reliability*, which we define as the *probability of delivering the required service (i.e., meeting the minimum performance requirements corresponding to the risk modification functionality) under specified operating conditions*. This is equivalent to the traditional notion of reliability, related to the probability of failure [10]; therefore we can adopt its formal definition (not given here). For our purposes, a barrier failure or *breach* constitutes those conditions when there is a deviation from the expected service, e.g., the barrier is unavailable when required (failure on demand), total loss of service, etc.

2) *Integrity*: Considering that not all failures lead to a safety related outcome (e.g., when a barrier fails safely), we define barrier *integrity* to be the *probability that it is not breached in a dangerous manner*. Thus, barrier reliability is related to *all* barrier failures whereas integrity considers only the *dangerous* failures.<sup>5</sup> Conservatively—i.e., when all barrier breaches are safety critical—reliability and integrity have the same value. The SORA concept of correction factor can now be seen as the logarithm of barrier integrity to the base ten.

3) *Assurance*: We can characterize assurance in a particular system property as *justified confidence* in that property. For measurable properties, such as reliability and integrity, we can relate confidence to the *uncertainty* in the measurement or estimation, which we can formally model using a probability distribution, through its parameters and moments [11]. For our purposes, assurance as it relates to barriers can be captured by (and is inversely proportional to) the uncertainty in barrier integrity. Thus, the greater the uncertainty in an integrity estimate, the lower the confidence and the lesser the assurance.

Based on this, we can give a mathematical interpretation to the SORA concepts of *integrity level* and *assurance level* as follows: partition the admissible range of quantitative values for integrity and assurance into disjoint intervals. Then, define a mapping from those intervals to the corresponding SORA measurement scale, i.e.,  $\{low, medium, high\}$ . To establish an appropriate mapping, we note that domain-specific input is required on the interval thresholds or bounds, which could be determined and validated by consortia such as JARUS.

<sup>5</sup>Although reliability affects safety, in general, a reliable system need not always be safe. However, since a barrier is a part of the safety system its correct or reliable operation is itself a safety requirement.

4) *Robustness*: The SORA concept of robustness is somewhat non-standard, being characterized as a combination of integrity and assurance. Since our formalization for integrity can be evaluated as a probability distribution that also quantifies assurance (as the uncertainty in the integrity, expressed by the distribution parameters and moments) the SORA notion of barrier robustness is implicitly captured.

SORA also includes a concept of *robustness level*, characterized as a combination of barrier integrity level and assurance level. It is straightforward to give a quantitative semantics to robustness level based on a formalization of assurance and integrity levels: a probability distribution of integrity, with the distribution parameters and moments determining the assurance, and the mapping to integrity and assurance level—given as in the preceding discussion—establishing the corresponding robustness level.

In general terms, however, (barrier) robustness is concerned with continued (barrier) operation to provide all, or a reasonable portion, of the expected functionality under off-nominal conditions. This is in contrast to reliability, where the operating conditions considered are typically nominal. In fact, to measure this traditional notion of robustness and rigorously capture its relation to risk modification, we must not only consider the specific off-nominal conditions where continued barrier operation is required, but also examine how the particular responses provided impact safety risk. Although we do not consider this further in the paper, as we will see later (Section V-A), our basis for risk assessment accounts for escalation factors, which can be seen as modeling, in part, specific off-nominal conditions that can compromise a barrier.

## B. Foundations for Risk Assessment

Recall that BTDs model (operational) risk scenarios that a given hazard induces, and the associated mitigation measures. Typically, we develop a BTD around a single top event of a hazard and, in general, a CONOPS will involve multiple hazards, each with possibly more than one top event. Thus, in practice, we can reasonably expect to develop a plurality of BTDs to capture the appropriate safety-related scenarios. Our previous work has extended BTDs, introducing and formalizing the concept of *safety architecture* (SA) [3]. Although the details are not in scope for this paper, in brief, an SA can be seen as a composition of related, mutually consistent BTDs for the identified hazards.

As such, an SA is relevant for SORA since it provides *i*) consistency across various operating scenarios, accounting for shared barriers and events, including threats and consequences, *ii*) a framework to view the full scope of safety in terms of the *safety system*—i.e., the overall collection of barriers—and how it contributes to SRM, *iii*) a more accurate and comprehensive context for risk assessment, and *iv*) the means to derive a rigorous foundation for risk assessment.

We now derive a simple, mathematical *risk model* based on BTDs, and extend it to the concept of SA. Using barrier integrities and (initial) threat probabilities, it focuses on computing the reduction in the probability of the consequences

and top events, as the foundation to establish whether the risk posed has been reduced to an acceptable level. Note that the model does not change the initial value of severity for the worst-case consequence(s) when determining the initial risk level (IRL) and the residual risk level (RRL).

Consider a BTD with a single threat  $T$ , top event  $\mathbf{T}$ , and single consequence  $C$ . The  $m$  prevention barriers on the incoming path  $\mathcal{I}$  from  $T$  to  $\mathbf{T}$ , are  $P_1, P_2, \dots, P_m$ . Similarly, let the  $q$  recovery barriers on the outgoing path  $\mathcal{O}$  from  $\mathbf{T}$  to  $C$  be  $R_1, R_2, \dots, R_q$ . We will abuse notation, letting  $T$ ,  $\mathbf{T}$ , and  $C$  be synonymous, respectively, with the Boolean random variables (RVs) modeling the threat, the top event, and the consequence respectively. Likewise,  $P_i$  and  $R_j$  are Boolean RVs modeling the prevention and recovery barriers, while the Boolean RVs  $\mathcal{I}$  and  $\mathcal{O}$  model the incoming and outgoing paths respectively, i.e., they model whether or not all the events on the path have occurred.

Except where otherwise stated, we use the convention  $\Pr(X)$  to mean the probability distribution over all states of the RV  $X$ , and  $\Pr(x)$  to mean the probability that  $X$  is in the state  $x$ . Since all the RVs in our case are Boolean, they can model states such as  $\{\text{operational}, \text{failed}\}$  (for barriers), and  $\{\text{occurs}, \text{does not occur}\}$  (in the case of threats, top events, consequences, and escalation factors).

We now analytically derive the expressions to compute risk (probability) reduction. Let  $\Pr(t)$  represent the initial probability that the threat  $T$  occurs, i.e.,  $\Pr(T = \text{true}) = \Pr(t)$ . Similarly,  $\Pr(\tau)$  is the marginal (i.e., unconditional) probability for the occurrence of the top event, and  $\Pr(c)$  the marginal consequence occurrence probability. We let  $\Pr(p_i)$ , equivalently  $\Pr(r_j)$ , be the unconditional barrier *fragility*—denoting the opposite of barrier integrity—defined such that integrity of the (prevention) barrier  $P_i$ , is  $1 - \Pr(p_i)$ .

The probability of the incoming path,  $\Pr(\mathcal{I})$ , is given as the joint probability of all the events on that path,  $\Pr(T, P_1, \dots, P_m, \mathbf{T})$ . To evaluate and simplify this expression, we assume that *i*) the barriers are mutually independent, and *ii*) barrier breaches are mutually independent of the threats or top events. Also, from BTD semantics, we observe that the top event occurs if the threat occurs and all barriers are breached. Based on these, and by applying the chain rule of probability, it can be shown that the incoming path probability is the product of the initial threat probability and the fragility of all the barriers on the path. That is,  $\Pr(\mathcal{I} = \text{true}) = \Pr(t) \prod_{i=1}^m \Pr(p_i)$ . The probability of the outgoing path  $\mathcal{O}$  can be similarly derived.

In the general case, i.e., as shown in Fig. 3, the top event occurs if all the events occur on any given path. That is, the top event occurs when any path occurs so that  $\Pr(\tau)$  is evaluated as the probability of the disjunction of the RVs for the  $w$  paths. That is,  $\Pr(\mathbf{T}) = \Pr(\bigcup_{i=1}^w \mathcal{I}_i)$ , which we can compute using the *inclusion-exclusion principle*. Since we require threats to be independent, the joint probability of a combination of threats occurring when evaluating this expression is simply the product of the individual threat probabilities. Also, a BTD has exactly one path between a top event and any

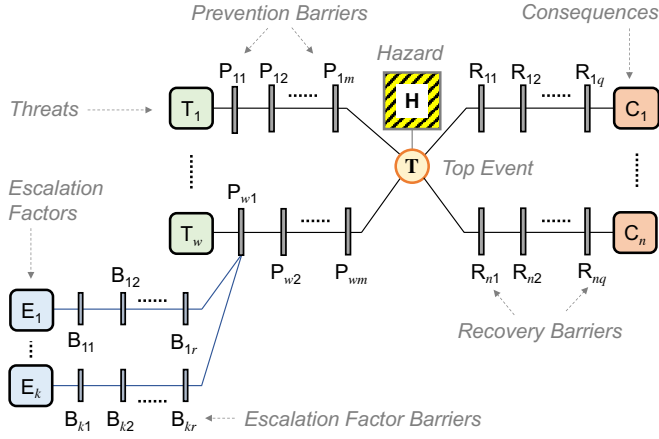


Fig. 3. Generic representation of a BTM, for a hazard  $H$  and top event  $T$ , showing its threats and consequences, a collection of prevention and recovery barriers, along with the escalation factors (EFs) and the escalation factor barriers (EFBs) for a specific prevention barrier.

given consequence  $C_i$ . Hence, the residual probability of that consequence (in the given BTM) is the probability of that path, given by generalizing the expression for  $\Pr(O = \text{true})$ , as  $\Pr(c_i) = \Pr(\tau) \prod_{j=1}^q \Pr(r_{ij})$ .

However, in an SA there can be several distinct paths to that consequence, each from a different top event. If there are  $n$  such paths then for a consequence  $C$ , we compute  $\Pr(C)$  as the probability of the disjunction of the RVs for the  $n$  paths. That is,  $\Pr(C) = \Pr(\bigcup_{i=1}^n S_i)$ , which we can compute using the inclusion-exclusion principle, as for the top event.

The details of extending our formal risk model to the entire SA, as well as uncertainty analysis, are out of scope for this paper. For uncertainty modeling and propagation, we can use the techniques developed in [11].

The preceding formalization addresses each of the SORA inconsistencies and open issues identified in Section III, as given in the order below. In particular, it:

i) explicitly captures the relation between unmitigated consequence risk and the contributing risk factors. Specifically, unmitigated consequence risk can be modeled as a BTM without barriers, whose threats or top events are the applicable risk factors. Then, by varying the latter and their associated parameters, e.g., threat probability, we can construct different scenarios, each with an assessment of initial unmitigated consequence risk. By grouping such scenarios into equivalence classes, we can give a formal basis for GRCs, tightly coupling them to the relevant risk factors.

ii) exactly models how consequence risk is modified by (harm) barrier correction factors, i.e., by formalizing both barrier correction factors (see Section IV-A) and the contribution of barrier (integrity) in estimating consequence probability and, thereby, consequence risk.

iii) explains how barrier robustness is related to barrier risk modification and consequence risk, i.e., by formalizing barrier robustness as a probability distribution for barrier integrity (see Section IV-A).

### C. Illustrative Example

We illustrate the applicability of our formal basis with a simple example extracted from a safety case we authored to support the NASA UAS Traffic Management (UTM) project. The intent is to show how a tool-supported implementation of our simple formal risk model can assist SORA.

The CONOPS on which our example is based involved conducting BVLOS flight on defined paths, with small UAs within an operating range (OR), a pre-defined volume of airspace that encloses, for the most part, sparsely populated and minimally built-up areas on the surface. The air traffic within and outside the OR includes conventionally piloted aircraft (CPA).

Fig. 4 shows one fragment from the different BTMs created to assess how midair collision (MAC) consequence risk would be reduced by deploying the barriers shown with the integrity values as indicated. The integrities have been given as nearest order of magnitude estimates, based upon available operational failure data, simulation data, and manufacturer specifications, where available, or from conservative assumptions as appropriate. Specifically, we use *ground-based surveillance* (with integrity 0.99), (a suite of) *avoidance maneuvers* (0.99), *emergency procedures* (0.5), an *independent flight abort mechanism* (0.999), also relying on *piloting actions* (0.9). From the SORA standpoint, these represent tactical mitigations.

Fig. 4 also shows how particular controls specialize the barriers, indicating the exact mitigation functionality used. Note that these correspond to the SORA recommended tactical mitigations. More generally, however, we can flexibly include a variety of barriers and controls, beyond those given in the SORA, whilst specifying their integrity. Based on the initial likelihood (probability) of the threat as shown, the risk assessment indicates that the top event (airborne conflict from a loss of separation) has a *medium* RRL, as does the MAC consequence, reduced from a *high* IRL. The risk levels and initial threat probability shown are based on mapping both the (specified and computed) event probabilities and severities to the likelihood and severity classes of a *risk classification and acceptance matrix*, e.g., as in [12].

Note that Fig. 4 represents the risk assessment of a specific scenario as modeled by a single BTM (shown here with a single threat and consequence). Our assurance case tool, Advocate [4], implements the formal risk model as extended to an SA, to expand the scope of such scenario-specific risk assessment and consolidate it over all constituent scenarios.

Although, presently, it uses point values for both barrier integrity and the initial threat probability to give point value results. Advocate automatically reconciles barrier repetition on a path between successive events, which occurs when different controls of the same barrier are used—e.g., as shown for the *ground-based surveillance* barrier in Fig. 4—also implementing well-formedness rules to preclude repetition of the same controls on a path.

The implementation computes the RRLs—in particular the residual probability of the consequence(s) and top event(s)—whereas the user specifies the relevant initial probabilities (and

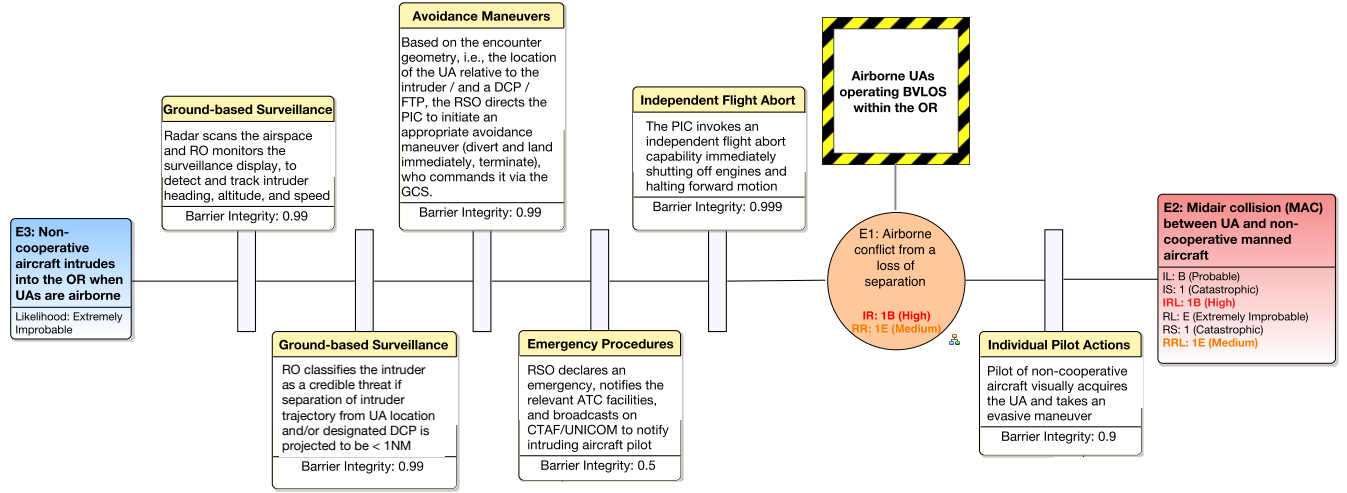


Fig. 4. Fragment of a BTd for a UAS CONOPS, showing a single threat, top event and consequence, along with the applicable prevention and recovery barriers and their respective integrities. The threat node also indicates the initial probability of occurrence, while the top event and the consequence show the initial and residual risk levels.

thereby also the IRLs). In principle, however, we can use the same approach to compute both the initial and residual probability for the events under consideration, by selecting the appropriate barriers to include in the analysis. Additionally, the implementation propagates the severity of the worst-case consequence to the precursor events, via the paths to the consequence.

## V. EXTENSIONS AND ENHANCEMENTS

### A. Formalizing Barrier Integrity

Based on Fig. 3 and our formal risk model (Section IV-B), an intuitive and straightforward extension is to provide a formal basis for determining barrier integrity, based on the identified EFs and EFBs. We note that SORA does not contain these concepts and, therefore, gives no recommendations on how EFs and EFBs should be considered within the process. Strictly speaking, although they are more relevant in barrier design and assurance rather than operational risk assessment, they play a critical role in operational safety, e.g., if an EF in one BTd can be a threat in a different BTd.

As such, in our opinion, considering this additional level of analysis can be useful and, effectively, we model barriers as a *second tier* BTd (see Fig. 3). Here, the top event is the barrier breach event, the threats are the identified EFs, and the prevention barriers are the EFBs respectively. Thus, if  $E$  is an EF to a barrier  $P_i$  that is managed using the EFBs  $B_1, B_2, \dots, B_r$ , then we can suitably modify the expression for  $\Pr(\mathcal{I})$ , as given in Section IV-B, to compute the barrier fragility as  $\Pr(p_i) = \Pr(e) \prod_{i=1}^r \Pr(b_i)$ . Moreover, if there are  $k$  escalation factors, i.e.,  $k$  paths  $\mathcal{E}_j$ , we can appropriately alter the expression for  $\Pr(\mathbf{T})$ , as developed in Section IV-B, to determine barrier fragility as,  $\Pr(P_i) = \Pr\left(\bigcup_{j=1}^k \mathcal{E}_j\right)$ .

### B. Implementation using Bayesian Networks

The implementation of the risk assessment foundations presented in Section IV-B and Section V-A, is limited in its

capacity to address practically relevant enhancements, such as multi-state discrete RVs, continuous RVs, distributions, modifications to the assumptions made in the analysis (e.g., of mutual independence amongst the barriers/controls), etc.

To address these gaps, we propose encoding the semantics of risk modification (as described by the scenarios that the BTds and SA capture) using Bayesian networks (BNs). Our motivation is their provision of a flexible probabilistic framework that affords efficient algorithms for reasoning under uncertainty, considering discrete and continuous RVs. Another key advantage is the specification of *prior* probabilities for the risk model parameters when there is insufficient information, and to update the priors, e.g., using operational data.

1) *Bayesian Networks*: A BN is a directed acyclic graph comprising nodes and edges, e.g., as shown in Fig. 5. Nodes represent (discrete, or continuous) RVs, the states of each of which are mutually exclusive. The directed edges capture dependencies between the RVs, and can be considered to model causal relationships. Thus, nodes linked by an edge are not independent and the link target, i.e., the child node, shows the dependent RV. We refer to nodes with no incoming edges as *root* nodes, nodes with both incoming and outgoing edges as *intermediate* nodes, and nodes with no outgoing edges as *leaf* nodes.

Associated with each RV in the BN is a (prior) *conditional probability table* (CPT) (for discrete RVs) or a *conditional probability distribution* (CPD) (for continuous RVs). For discrete RVs, the CPT specifies the conditional probability for each state of the RV given every possible combination of the states of its parents. If an RV has no parents, the CPT is the same as the unconditional probability. When all CPTs for all RVs of a BN have been specified, it effectively gives the *joint distribution* over the RVs, and the BN is completely specified. Then, we can compute the unconditional probability for any non-root RV by *marginalization* [13].

Moreover, using Bayes' theorem, we can query the BN to



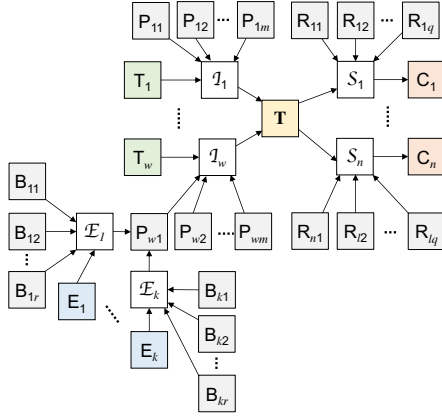


Fig. 5. BN representation of the BTD of Fig. 3.

determine the probability of a specific state of a specific RV, as well as the joint probability of any combination of RVs given *evidence*, i.e., observations on the values of other RV in the network.

2) *Implementation of the Risk Model*: Our intuition is to use the structure of a BTD (more generally, the SA) and its semantics, to derive the topology of the BN. For example, we can use root BN nodes to model threats, barriers, EFs, and EFBs, intermediate BN nodes to capture top events, and leaf BN nodes to model the various consequences.

Fig. 5 shows a BN representation for the generic BTD of Fig. 3, the nodes of which are Boolean RVs. To show the correspondence between the two, the identifiers of the RVs in BN of Fig. 5 have been made synonymous with those of the BTD elements of Fig. 3. The RVs  $\mathcal{I}_1 \dots \mathcal{I}_w$ ,  $\mathcal{S}_1 \dots \mathcal{S}_n$ , and  $\mathcal{E}_1 \dots \mathcal{E}_k$ , model (the state of) specific paths in the BTD, i.e., *incoming* paths from the threats to the top event, *outgoing* paths from the top event to the consequences, and *escalation* paths from the EFs to a barrier, respectively.

To complete the BN specification, we must also give the CPTs for the RVs. Root node CPTs are, trivially, the corresponding unconditional probabilities. For instance, for the RVs corresponding to the threats  $T_i$ , the CPTs are simply the initial probabilities of the threats  $\Pr(T_i)$ . For the intermediate and leaf nodes in the BN, the idea is that *a*) by applying the chain rule, *b*) using the assumptions made regarding independence relations amongst the barriers (and threats), and *c*) from the semantics of the BTDs, we can rearrange and reuse the mathematical expressions developed in Section IV-B and Section V-A as appropriate, to derive the relevant CPTs. For example, in Fig. 5, the CPT for the RV of the top event,  $\mathbf{T}$ , is  $\Pr(\mathbf{T}|\mathcal{I}_1, \dots, \mathcal{I}_w)$ , which is exactly given by the expression for  $\Pr(\mathbf{T})$  as in Section IV-B. In fact, this CPT is the truth table for a Boolean OR of the RVs  $\mathbf{T}, \mathcal{P}_1, \dots, \mathcal{P}_w$ . The CPTs for the remaining intermediate and leaf nodes can be derived in a similar way although, due to space constraints, we do not show it here.

The BN for the SA composes the various BNs that correspond to the constituent BTDs of the SA. The details of the composition are not in scope for this paper, and the key

point is that there can be additional parents for each threat  $T_i$  or consequence  $C_i$ . Those are the respective path RVs of the *other* BTDs in the SA. Thus, in the BN of the SA, we must update the relevant CPT. For example, for a consequence  $C_i$  we must modify the CPT to  $\Pr(C_i|\mathcal{S}_i, \mathcal{S}_j)$ , where  $j$  is an index over the  $b$  BTDs containing the consequence modeled by the RV  $C_i$ . By properly modifying the expression for  $\Pr(C)$  given in Section IV-B, we can exactly specify the required CPT as  $\Pr(C_i|\mathcal{S}_i, \mathcal{S}_j) = \Pr(\bigcup_{j=1}^b \mathcal{S}_j \cup \mathcal{S}_i)$ .

Quantifying risk reduction then amounts to querying the BN of the SA, i.e., evaluating the marginal distribution of the specific RVs that model the consequences,  $C_1, \dots, C_n$ .

## VI. DISCUSSION AND CONCLUSIONS

We have summarized the SORA methodology, highlighting a number of inconsistencies that motivate the benefit of a formal underpinning. We have advanced a well-founded, albeit preliminary, approach for quantifying risk reduction also explaining how our formalization addresses the identified concerns. Our approach is supported by mathematically-based notions of the relevant concepts that are quantitatively verifiable, e.g., barrier integrity can be corroborated through measurement, simulation, and operational monitoring, etc. We have also illustrated the applicability of our approach through a simple example drawn from a UAS safety case, following which we have described a preliminary extension (to determine barrier integrity), as well as a BN-based enhancement.

The latter offers key opportunities to further enhance risk assessment. For example, from the standpoint of designing the safety system, starting from a safety target we can allocate risk—in particular, its probability component—across the various barriers known to be independent. A simplistic approach is to equally allocate risk across all barriers. A more well-founded approach could be to use sensitivity analysis [14], to evaluate how barriers and risk factors contribute to overall safety risk, and use the results to guide risk allocation. From an assessment standpoint, such analysis can also provide insight into the collection of barriers that have an appreciable risk reduction impact, thereby guiding the selection of barriers appropriate for the risk posed.

BNs also offer a framework in which to extend our risk model, by including other models used in aviation safety that have a probabilistic formulation [15], [16]. These could be either directly used as input, e.g., to specify appropriate CPTs, or transformed to be included as part of our overall risk model. In either case, a key advantage that we perceive is accessibility to measurable model parameters that, in turn, could update the risk assessment based on operational monitoring.

Others have also investigated applying BNs for UAS risk assessment [17]. The development of a principled basis to specify the BN, based on the semantics of the safety scenarios modeled by the BTDs and the SA, differentiates our approach from this prior work. Along those lines, the mapping from BTDs to BNs also has been considered previously [18], although the approach there is to transform BTDs, first, into fault trees and event trees, and then translate those to a BN.

Additionally, there is no notion of an SA as a composition of BTDs that is then transformed to a BN.

A number of challenges need to be addressed to further enhance the practical utility of our approach. First, assumptions made for risk analysis require careful consideration and justification, e.g., barrier independence. Additionally, integrity quantification can be problematic when barriers involve software, or human operators and procedures, requiring careful justification when quantified. To address the former, we can leverage our prior work on *structured assurance arguments* [19], to substantiate specific (assurance) claims with concrete evidence, through a chain of reasoning. For the latter, we envision two alternatives: *i)* conduct the risk assessment considering only those barriers that can be justifiably quantified, to gauge the amount of residual risk to be addressed by human factors and software. Then, using structured arguments, give rationale for any claims of *fitness for purpose* and risk reduction, supported by concrete verifiable evidence; or, *ii)* assume *non-informative objective priors* on barrier integrities (e.g., barrier functionality is equally likely to be effective or ineffective), to gauge the amount of additional risk reduction required from the remainder of the barriers whose integrity can be measured.

Moreover, the assumption of independence among barriers, and between barriers and threats, may not always hold, e.g., when barriers are coupled through a common cause, or when there are other interdependencies. Here, we must update the models with *conditional* probabilities. In the BN-based implementation, we can capture dependencies by introducing appropriately directed edges between the RVs involved, and also updating the CPTs of the dependent RVs.

Finally, our formal basis and its implementation can be improved along several lines. A key enhancement is to work at the level of controls in the risk assessment, which will require considering barrier configurations and organization. While this additional level of detail may not be justified for lower-risk operations, the need for higher assurance and improved accuracy of the risk analysis in high-risk operations may warrant this finer granularity. We also plan to include uncertainty analysis and propagation into the underlying formalization and its implementation, to rigorously address the concerns of uncertainty in the assessment.

The SORA largely adopts qualitative risk assessment to provide flexibility both to the NAAs in employing guidance relevant to their state, and to UAS operators to demonstrate that their safety measures are fit for purpose, relative to the risk associated with their operations. The preference for a qualitative approach cites challenges in data collection, and in managing various associated uncertainties. However, a qualitative approach may be problematic for operations in more complex, higher-risk environments. This paper underscores the general observation that the results of quantitative analysis under uncertainty require careful treatment. Nevertheless, it also advocates that a qualitative analysis that is informed by, and founded on, the rigor of quantification may be preferable. Additionally, quantification can be advantageous, complementing the SORA methodology by facilitating measurable jus-

tification of qualitative assumptions, precluding (deceptively convincing) inconsistencies, whilst retaining its flexibility.

## ACKNOWLEDGMENT

This work was supported in part by the SWS project, and the UTM project, under the Airspace Operations and Safety Program of NASA ARMD. We also thank Terry Martin (QUT, Australia) for early discussions and feedback on some of the ideas developed in this paper.

## REFERENCES

- [1] Joint Authorities for Rulemaking on Unmanned Systems, "JARUS guidelines on Specific Operations Risk Assessment (SORA)," Final Public Release v. 1.0, Jun. 2017. [Online]. Available: <http://jarus-rpas.org/content/jar-doc-06-sora-package>
- [2] European Aviation Safety Agency, "Concept of Operations for Drones. A Risk Based Approach to Regulation of Unmanned Aircraft," EASA Brochure, May 2015.
- [3] E. Denney, G. Pai, and I. Whiteside, "Modeling the Safety Architecture of UAS Flight Operations," in *Computer Safety, Reliability, and Security (SAFECOMP 2017)*, LNCS vol. 10488, Sep. 2017, pp. 162–178.
- [4] E. Denney, G. Pai, and I. Whiteside, "Model-driven development of safety architectures," in *2017 ACM/IEEE 20th Intl. Conf. Model Driven Engineering Languages and Systems (MODELS 2017)*, Sep. 2017, pp. 156–166.
- [5] E. Denney and G. Pai, "Safety considerations for UAS ground-based detect and avoid," in *2016 IEEE/AIAA 35th Digital Avionics Systems Conf. (DASC 2016)*, Sep. 2016, pp. 1–10.
- [6] M. Johnson, J. Jung, J. Rios, J. Mercer, J. Homola, T. Prevot, D. Mulfinger, and P. Kopardekar, "Flight Test Evaluation of an Unmanned Aircraft System Traffic Management (UTM) Concept for Multiple Beyond-Visual-Line-of-Sight Operations," in *12th USA/Europe Air Traffic Management Research and Development Seminar (ATM2017)*, Jun. 2017.
- [7] R. Clothier, E. Denney, and G. Pai, "Making a Risk Informed Safety Case for Small Unmanned Aircraft System Operations," in *17th AIAA Aviation Technology, Integration, and Operations Conf., AIAA Aviation Forum*, AIAA 2017-3275, Jun. 2017.
- [8] E. Denney and G. Pai, "Tool support for assurance case development," *Journal of Automated Software Engineering*, vol. 25, no. 3, pp. 435–499, Sep. 2018.
- [9] F. S. Roberts, *Measurement Theory With Applications to Decision making, Utility, and the Social Sciences*, Encyclopedia of Mathematics and its Applications, vol. 7, Cambridge Univ. Press, Mar. 2009.
- [10] K. Trivedi and A. Bobbio, *Reliability and Availability Engineering. Modeling, Analysis, and Applications*. Cambridge Univ. Press, 2017.
- [11] L. Yin, M. A. J. Smith, and K. S. Trivedi, "Uncertainty Analysis in Reliability Modeling," in *Proceedings of the 2001 Annual Reliability and Maintainability Symp. (RAMS 2001)*, 2001, pp. 229–234.
- [12] FAA Air Traffic Organization, *Safety Management System Manual version 4.0*, Federal Aviation Administration, May 2014.
- [13] D. Barber, *Bayesian Reasoning and Machine Learning*. Cambridge Univ. Press, 2012.
- [14] H. Chan and A. Darwiche, "Sensitivity Analysis in Bayesian Networks: From Single to Multiple Parameters," in *Proc. 20th Conf. on Uncertainty in Artificial Intelligence (UAI 2004)*, 2004, pp. 67–75.
- [15] R. E. Weibel and R. J. Hansman, "Safety Considerations for Operation of Different Classes of UAVs in the NAS," in *AIAA 3rd "Unmanned Unlimited" Technical Conf., Workshop and Exhibit, Infotech@Aerospace Conferences*, AIAA 2004-6421, Sep. 2004.
- [16] A. McFadyen, T. Martin, and T. Perez, "Low-level collision risk modelling for unmanned aircraft integration and management," in *Proc. 2018 IEEE Aerospace Conf.*, March 2018, pp. 1–10.
- [17] E. Ancel, F. Capristan, J. Foster, and R. Condotta, "Real-time Risk Assessment Framework for Unmanned Aircraft System (UAS) Traffic Management (UTM)," in *17th AIAA Aviation Technology, Integration, and Operations Conf., AIAA Aviation Forum*, AIAA 2017-3273, 2017.
- [18] N. Khakzad, F. Khan, and P. Amyotte, "Dynamic Safety Analysis of Process Systems by mapping Bow-Tie into Bayesian network," *Process Safety and Environmental Protection*, vol. 91, no. 1, pp. 46 – 53, 2013.
- [19] E. Denney and G. Pai, "A methodology for the development of assurance arguments for unmanned aircraft systems," in *33rd Intl. System Safety Conf. (ISSC 2015)*, Aug. 2015.